

Remarks

Applicants have reviewed this Application in light of the Office Action sent electronically 7 May 2007. Applicants have made clarifying amendments to Claims 1, 5-6, 9-11, 13-15, 17-19, and 21-24 and added new Claims 25-26. Applicants respectfully request reconsideration and allowance of all pending claims.

Information Disclosure Statement

The Examiner has not considered References C-D, G-H, and J in the PTO-1449 accompanying the Information Disclosure Statement (IDS) filed 13 November 2006. The Examiner asserts, "for a reference, listed in Non Patent Literature Documents, it must submit a copy and must be identified by publisher, author (if any), title, relevant pages of the publication, date, and place of publication. If a reference that does not meet the requirement under 1.98(b), it will result in being strikethrough."

The IDS filed 13 November 2006 included legible copies of References C-D, G-H, and J and their International Search Reports, which provide most if not all the information sought by 37 C.F.R. §1.98(b)(5). Applicants respectfully submit that any information missing from the listing of these References in the accompanying PTO-1449 does not in any way hinder the Examiner's ability to consider them. Nonetheless, to expedite consideration of these References, Applicants have submitted along with this Response a new IDS with a new PTO-1449 listing these References with additional identification information. Applicants respectfully request the Examiner to consider these References.

The Claims are Definite

The Examiner rejects Claims 1-24 under 35 U.S.C. § 112 para. 2 as being indefinite. Specifically, the Examiner asserts that *an integrated fabric*, as independent Claim 1 recites, is ambiguous. The Examiner applies the same rejection to independent Claims 9 and 17. Applicants respectfully disagree with the Examiner.

According to the M.P.E.P., “breadth of a claim is not to be equated with indefiniteness.” M.P.E.P. ch. 2173.04 (Rev. 3, August 2005). “If the claims, read in light of the specification, reasonably apprise those skilled in the art . . . of the utilization and scope of the invention, and if the language is as precise as the subject matter permits, [35 U.S.C. § 112 para. 2] demands no more.” M.P.E.P. ch. 2173.05(a)(II) (Rev. 3, August 2005). Similarly, if a claim is directed to patentable subject matter, the claim is allowable if it “define[s] the patentable subject matter with a reasonable degree of particularity and distinctness.” M.P.E.P. ch. 2173.02 (Rev. 3, August 2005) (emphasis in original). “Some latitude in the manner of expression and the aptness of terms should be permitted even though the claim language is not as precise as the examiner might desire.” *Id.* The M.P.E.P. also states:

Definiteness of claim language must be analyzed, not in a vacuum, but in light of:

- (A) The content of the particular application disclosure;
- (B) The teachings of the prior art; and
- (C) The claim interpretation that would be given by one possessing the ordinary level of skill in the pertinent art at the time the invention was made.

Id.

With respect to *an integrated fabric*, as independent Claim 1 recites, the Specification provides the following descriptions of example nodes:

FIGURES 3A-C illustrate various embodiments of individual nodes 115 in grid 110. In the illustrated, but example, embodiments, nodes 115 are represented by blades 315. Blade 315 comprises any computing device in any orientation operable to process all or a portion, such as a thread or process, of job 150. For example, blade 315 may be a standard Xeon64™ motherboard, a standard PCI-Express Opteron™ motherboard, or any other suitable computing card.

Blade 315 is an integrated fabric architecture that distributes the fabric switching components uniformly across nodes 115 in grid 110, thereby possibly reducing or eliminating any centralized switching function, increasing the fault tolerance, and allowing message passing in parallel. More specifically, blade 315 includes an integrated switch 345. Switch 345 includes any number of ports that may allow for different topologies. For example, switch 345 may be an eight-port switch that enables a tighter three-dimensional mesh or 3D Torus topology. These eight ports include two “X” connections for linking to neighbor nodes 115 along an X-axis, two “Y” connections for linking to neighbor nodes 115 along a Y-axis, two “Z”

connections for linking to neighbor nodes 115 along a Z-axis, and two connections for linking to management node 105. In one embodiment, switch 345 may be a standard eight port Infiniband-4x switch IC, thereby easily providing built-in fabric switching. Switch 345 may also comprise a twenty-four port switch that allows for multidimensional topologies, such as a 4-D Torus, or other non-traditional topologies of greater than three dimensions. Moreover, nodes 115 may further interconnected along a diagonal axis, thereby reducing jumps or hops of communications between relatively distant nodes 115. For example, a first node 115 may be connected with a second node 115 that physically resides along a northeasterly axis several three dimensional "jumps" away.

FIGURE 3A illustrates a blade 315 that, at a high level, includes at least two processors 320a and 320b, local or remote memory 340, and integrated switch (or fabric) 345. Processor 320 executes instructions and manipulates data to perform the operations of blade 315 such as, for example, a central processing unit (CPU). Reference to processor 320 is meant to include multiple processors 320 where applicable. In one embodiment, processor 320 may comprise a Xeon64 or Itanium™ processor or other similar processor or derivative thereof. For example, the Xeon64 processor may be a 3.4GHz chip with a 2MB Cache and HyperTreading. In this embodiment, the dual processor module may include a native PCI/Express that improves efficiency. Accordingly, processor 320 has efficient memory bandwidth and, typically, has the memory controller built into the processor chip.

Blade 315 may also include Northbridge 321, Southbridge 322, PCI channel 325, HCA 335, and memory 340. Northbridge 321 communicates with processor 320 and controls communications with memory 340, a PCI bus, Level 2 cache, and any other related components. In one embodiment, Northbridge 321 communicates with processor 320 using the frontside bus (FSB). Southbridge 322 manages many of the input/output (I/O) functions of blade 315. In another embodiment, blade 315 may implement the Intel Hub Architecture (IHA™), which includes a Graphics and AGP Memory Controller Hub (GMCH) and an I/O Controller Hub (ICH).

PCI channel 325 comprises any high-speed, low latency link designed to increase the communication speed between integrated components. This helps reduce the number of buses in blade 315, which can reduce system bottlenecks. HCA 335 comprises any component providing channel-based I/O within server 102. Each HCA 335 may provide a total bandwidth of 2.65GB/sec, thereby allowing 1.85GB/sec per PE to switch 345 and 800MB/sec per PE to I/O such as, for example, BIOS (Basic Input/Output System), an Ethernet management interface, and others. This further allows the total switch 345 bandwidth to be 3.7GB/sec for 13.6Gigaflops/sec peak or 0.27Bytes/Flop I/O rate is 50MB/sec per Gigaflop.

Memory 340 includes any memory or database module and may take the form of volatile or non-volatile memory including, without limitation, magnetic media, optical media, flash memory, random access memory (RAM), read-only memory (ROM), removable media, or any other suitable local or remote memory component. In the illustrated embodiment, memory 340 is comprised of 8 GB of dual double data rate (DDR) memory

components operating at least 6.4GB/s. Memory 340 may include any appropriate data for managing or executing HPC jobs 150 without departing from this disclosure.

FIGURE 3B illustrates a blade 315 that includes two processors 320a and 320b, memory 340, HyperTransport/ peripheral component interconnect (HT/PCI) bridges 330a and 330b, and two HCAs 335a and 335b.

Example blade 315 includes at least two processors 320. Processor 320 executes instructions and manipulates data to perform the operations of blade 315 such as, for example, a central processing unit (CPU). In the illustrated embodiment, processor 320 may comprise an Opteron processor or other similar processor or derivative. In this embodiment, the Opteron processor design supports the development of a well balanced building block for grid 110. Regardless, the dual processor module may provide four to five GigaFlop usable performance and the next generation technology helps solve memory bandwidth limitation. But blade 315 may more than two processors 320 without departing from the scope of this disclosure. Accordingly, processor 320 has efficient memory bandwidth and, typically, has the memory controller built into the processor chip. In this embodiment, each processor 320 has one or more HyperTransport™ (or other similar conduit type) links 325.

Generally, HT link 325 comprises any high-speed, low latency link designed to increase the communication speed between integrated components. This helps reduce the number of buses in blade 315, which can reduce system bottlenecks. HT link 325 supports processor to processor communications for cache coherent multiprocessor blades 315. Using HT links 325, up to eight processors 320 may be placed on blade 315. If utilized, HyperTransport may provide bandwidth of 6.4 GB/sec, 12.8, or more, thereby providing a better than forty-fold increase in data throughput over legacy PCI buses. Further HyperTransport technology may be compatible with legacy I/O standards, such as PCI, and other technologies, such as PCI-X.

Blade 315 further includes HT/PCI bridge 330 and HCA 335. PCI bridge 330 may be designed in compliance with PCI Local Bus Specification Revision 2.2 or 3.0 or PCI Express Base Specification 1.0a or any derivatives thereof. HCA 335 comprises any component providing channel-based I/O within server 102. In one embodiment, HCA 335 comprises an Infiniband HCA. InfiniBand channels are typically created by attaching host channel adapters and target channel adapters, which enable remote storage and network connectivity into an InfiniBand fabric, illustrated in more detail in FIGURE 3B. Hypertransport 325 to PCI-Express Bridge 330 and HCA 335 may create a full-duplex 2GB/sec I/O channel for each processor 320. In certain embodiments, this provides sufficient bandwidth to support processor-processor communications in distributed HPC environment 100. Further, this provides blade 315 with I/O performance nearly or substantially balanced with the performance of processors 320.

FIGURE 3C illustrates another embodiment of blade 315 including a daughter board. In this embodiment, the daughter board may support 3.2GB/sec or higher cache coherent interfaces. The daughter board is operable to include one or more Field Programmable Gate Arrays (FPGAs)

350. For example, the illustrated daughter board includes two FPGAs 350, represented by 350a and 350b, respectively. Generally, FPGA 350 provides blade 315 with non-standard interfaces, the ability to process custom algorithms, vector processors for signal, image, or encryption/decryption processing applications, and high bandwidth. For example, FPGA may supplement the ability of blade 315 by providing acceleration factors of ten to twenty times the performance of a general purpose processor for special functions such as, for example, low precision Fast Fourier Transform (FFT) and matrix arithmetic functions.

The preceding illustrations and accompanying descriptions provide exemplary diagrams for implementing various scaleable nodes 115 (illustrated as example blades 315). However, these figures are merely illustrative and system 100 contemplates using any suitable combination and arrangement of elements for implementing various scalability schemes. Although the present invention has been illustrated and described, in part, in regard to blade server 102, those of ordinary skill in the art will recognize that the teachings of the present invention may be applied to any clustered HPC server environment. Accordingly, such clustered servers 102 that incorporate the techniques described herein may be local or a distributed without departing from the scope of this disclosure. Thus, these servers 102 may include HPC modules (or nodes 115) incorporating any suitable combination and arrangement of elements for providing high performance computing power, while reducing I/O latency. Moreover, the operations of the various illustrated HPC modules may be combined and/or separated as appropriate. For example, grid 110 may include a plurality of substantially similar nodes 115 or various nodes 115 implementing differing hardware or fabric architecture.

(Figures 3A-3C; Specification, Page 18, Line 12, through Page 24, Line 12).

Applicants respectfully submit that, when analyzed in light of the above disclosure, as well as the teachings of the prior art and the interpretation that a person having ordinary skill in the art at the time of the invention would give independent Claim 1, *an integrated fabric*, as independent Claim 1 recites, is clearly definite under 35 U.S.C. §112 para. 2.

The Examiner further asserts that *determining an unallocated subset from a plurality of high performance computing nodes*, as independent Claim 1 recites, is “ambiguous because it mentally indicates an aspect rather physically points out a subject matter. It should be noted that before an allocation, every element is unallocated, therefore saying ‘determining an unallocated subset from a plurality of HPC nodes’ is a mere mental indication, and it does not know where an existed element according to the claims is allocated or unallocated.” The Examiner applies the same rejection to independent Claims 9

and 17. Applicants respectfully do not understand the rationale given for this rejection. Applicants respectfully submit that *determining an unallocated subset from a plurality of high performance computing nodes*, as independent Claim 1 recites, at least “reasonably apprises those skilled in the art of the utilization and scope of the invention,” at least “define[s] the patentable subject matter with a reasonable degree of particularity and distinctness,” and is therefore definite under 35 U.S.C. §112 para. 2.

Accordingly, Applicants respectfully request reconsideration and allowance of independent Claims 1, 9, and 17 and all their dependent claims.

Claims 9-24 Recite Patentable Subject Matter

The Examiner rejects Claims 9-16 under 35 U.S.C. § 101 as being directed to nonstatutory subject matter. Specifically, the Examiner asserts, “Claims 9-16 recites software. The claims do not meet the statutory requirement because the claims are programming per se.” To expedite issuance of a patent from this Application, Applicants have made clarifying amendments to independent Claim 9. Applicants respectfully request reconsideration and allowance of Claims 9-16.

The Examiner also rejects Claims 17-24 under 35 U.S.C. § 101 as being directed to nonstatutory subject matter. Specifically, the Examiner asserts, “Claims 17-24 recites a system for job management, where ‘system’ is generic. The claims are generally associated with an ‘HPC environment,’ where environment is generic.” To expedite issuance of a patent from this Application, Applicants have made clarifying amendments to independent Claim 17. Applicants respectfully request reconsideration and allowance of Claims 17-24.

Independent Claims 1, 9, and 17 are Allowable Over Keller

The Examiner rejects independent Claims 1, 9, and 17 under 35 U.S.C. § 102(a) as being anticipated by Hovestadt, Kao, Keller, and Streit, *Scheduling in HPC Resource Management Systems: Queuing vs. Planning*, PROCEEDINGS OF THE 9TH WORKSHOP ON JOB SCHEDULING STRATEGIES FOR PARALLEL PROCESSING (“Keller”).

Keller merely discloses a resource management system (RMS) for high performance computing (HPC) machines that closes a gap between the RMS and grid middleware. (Page 1). The RMS in *Keller* provides knowledge about start times of all requests in the system, making advanced reservations implicitly possible. (Page 2).

In contrast, independent Claim 1, as amended, recites:

A method comprising:
determining an unallocated subset from a plurality of high performance computing (HPC) nodes, each of the unallocated HPC nodes comprising an integrated fabric;
selecting an HPC job from a job queue; and
executing the selected job using at least a portion of the unallocated subset of nodes.

Independent Claims 9 and 17 are similar to independent Claim 1.

Keller fails to disclose, teach, or suggest each and every limitation of independent Claim 1. As an example, *Keller* fails to disclose, teach, or suggest ***determining an unallocated subset from a plurality of high performance computing (HPC) nodes, each of the unallocated HPC nodes comprising an integrated fabric***, as independent Claim 1 recites. Even assuming for the sake of argument that the HPC machines in *Keller* could be properly considered ***a plurality of high performance computing (HPC) nodes***, as independent Claim 1 recites, *Keller* would still fail to disclose, teach, or suggest any of the HPC machines in *Keller* ***comprising an integrated fabric***, as independent Claim 1 recites, much less ***each of the unallocated HPC machines in Keller comprising an integrated fabric***, as independent Claim 1 recites.

“To anticipate, every element and limitation of the claimed invention must be found in a single prior art reference, arranged as in the claim.” *Brown v. 3M*, 265 F.3d 1349, 1351 (Fed. Cir. 2001). “A claim is anticipated only if each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference.” *Verdegaal Bros. v. Union Oil Co. of California*, 814 F.2d 628, 631, 2 U.S.P.Q.2d 1051, 1053 (Fed. Cir. 1987); M.P.E.P. ch. 2131 (Rev. 3, Aug. 2005) (quoting *Verdegaal*, 814 F.2d at

631). Moreover, “[t]he identical invention must be shown in as complete detail as is contained in the patent claim.” *Richardson v. Suzuki Motor Co.*, 868 F.2d 1226, 1236, 9 U.S.P.Q.2d 1913, 1920 (Fed. Cir. 1989); M.P.E.P. ch. 2131 (Rev. 3, Aug. 2005) (quoting *Richardson*, 868 F.2d at 1236). Furthermore, “[t]he elements must be arranged as required by the claim.” M.P.E.P. ch. 2131 (Rev. 3, Aug. 2005) (citing *In re Bond*, 910 F.2d 831, 832, 15 U.S.P.Q.2d 1566, 1567 (Fed. Cir. 1990)). As shown above, *Keller* fails to disclose, either expressly or inherently, each and every limitation of independent Claim 1. Therefore, *Keller* does not anticipate independent Claim 1 under governing Federal Circuit case law and the M.P.E.P.

For at least the reasons above, Applicants respectfully request reconsideration and allowance of independent Claims 1, 9, and 17 and all their dependent claims.

Conclusion

For at least the foregoing reasons, Applicants respectfully request allowance of all pending claims.

If a telephone conference would advance prosecution of this Application, Applicants invite the Examiner to call Travis W. Thomas, Attorney for Applicants, at 214.953.6676.

Please charge \$1,050.00 for a three-month extension of time to Deposit Account No. 02-0384 of Baker Botts LLP. The Commissioner may charge any fee and credit any overpayment to Deposit Account No. 02-0384 of Baker Botts LLP.

Respectfully submitted,
BAKER BOTTS L.L.P.
Attorneys for Applicants

A handwritten signature in black ink, appearing to be 'T. T.' with a horizontal line extending from the second 'T'.

Travis W. Thomas
Reg. No. 48,667

Date: 7 October 2007

Correspondence Address:

Customer Number: 45507